

# Suspending Recursion in Causal-link Planning

**David E. Smith**

Rockwell Science Center  
444 High St.  
Palo Alto, CA 94301  
de2smith@rpal.rockwell.com  
(415) 325-7162

**Mark A. Peot**

Rockwell Science Center  
444 High St.  
Palo Alto, CA 94301  
peot@rpal.rockwell.com  
(415) 325-7143

## **Abstract**

We present a strategy for suspending recursive open conditions during planning. We also show conditions under which plans with suspended open conditions can be pruned. To make this suspension and pruning strategy efficient, we use an *operator graph* to analyze potential recursion before the planning process begins.

This approach covers a broader range of recursive problems than the approaches of Morris and Kambhampati, and is much more tractable than Kambhampati's approach. We give experimental results that indicate 1) significant improvement on recursive problems and 2) negligible overhead when applied to recursive and non-recursive problems alike.

**Keywords:** causal-link planning, recursion, search-control, open condition ordering.

# 1 Introduction

Recursion is prevalent in many planning domains. In fact, any time some of the operations are reversible, there is the possibility of recursion in the planning process. For transportation problems recursion occurs whenever it is possible to travel both directions along a road or airway, or whenever it is possible to both load and unload cargo. In Russell's tire-changing problem, there is an operator for opening the car trunk and another for closing it. Both operators are necessary to achieve the goals of changing the tire and putting all the tools away. Unfortunately, the presence of these two operators allows arbitrarily long sequences of Open-trunk, Close-trunk, Open-trunk, Close-trunk ...

Recursion problems like this can often be alleviated by smart search strategies. For example, a planner can rank partial plans according to the amount of recursion and prefer to work on those partial plans with less recursion. Once a partial plan is chosen for elaboration, the planner can prefer working on open conditions that are not recursive. While such control strategies are often useful, if the planning problem has no solution, recursive plans and open conditions will eventually be found and doggedly explored by the planner until the end of time (or memory).

Some planning systems avoid recursion by either pruning [5, 11] or restricting the expansion of recursive open conditions [15]. Feldman and Morris [3] have pointed out that these techniques are, in general, not admissible. (Sometimes they prune the only viable plans). In [3], Feldman and Morris prove certain conditions under which it is possible to stop expansion of recursive open conditions. In some cases, it is necessary for them to add filter preconditions to operators in order for the theorems to apply. It is not clear what fraction of recursion problems can, in fact, be handled by their method.

More recently, Kambhampati [7] has presented methods for pruning *non-minimal* partial plans during search. While the theory is general, there are two significant problems with its practical use:

1. The technique is computationally expensive for causal-link planners.
2. The conditions for pruning require that all remaining open-conditions in a partial plan temporally precede the root of the recursion (see [7] for a precise description). This means that a partial plan must be at an advanced stage of development before pruning is possible.

In this paper we also present admissible techniques for pruning recursion. In Section 2 we present a method that suspends and prunes repeating open conditions. As with the techniques of Feldman and Morris, and Kambhampati, this method applies to recursion where the open condition is repeated exactly. The technique we develop is relatively cheap and catches recursion early.

While the technique in Section 2 is applicable as described, there are ways to make it significantly more efficient and powerful. In Section 3 we discuss the use of an *operator graph* (introduced in [13]) for:

1. Quick detection of recursion during planning
2. Early pruning of plans with suspended open conditions


In Section 4 we consider a much broader class of recursion that we call *instance recursion*. In this type of recursion, the predicate repeats, but the variable bindings may be different each time. We develop a suspension and pruning technique that controls this type of recursion.

For purposes of this paper we will consider a simple SNLP style causal-link planner [8, 1]. To indicate the operator used in a particular plan step we will follow the step name by the operator expression, e.g.  $S_5 : \text{Go}(x, B)$ . With recursion, open conditions share the same predicate and variable bindings, so we will explicitly name open conditions and follow with the actual expression, e.g.  $O_3 : \text{At}(x)$ . The expression  $S_1 \xrightarrow{C} S_2$  will be used to denote the causal link from step  $S_1$  to  $S_2$  with the condition  $C$ .

Due to space limitations we provide only informal arguments for the theorems in this paper.

## 2 Exact Recursion

Consider the partial plan shown in Figure 1. In this plan, the open condition  $O:C$  appears as a subgoal of the condition  $C$  in the causal link  $S_j \xrightarrow{C} S_i$ . Intuitively it seems that we should be able to discard this partial plan. If there is some way of achieving  $O:C$  then it could presumably be used to achieve  $C$  directly, without the intervening steps from  $S_k$  to  $S_j$ . Unfortunately, this argument is not generally sound.




**Figure 1:** Simple example of exact recursion.


There are two circumstances where the steps between  $C$  and  $O:C$  are still needed:

1. Some step between  $S_k$  and  $S_j$  (inclusive) might be used for another purpose, i.e., some other open condition might link into one of these steps.
2. Some other step that clobbers  $C$  might be forced to come between  $S_k$  and  $S_j$ . In other words, a threat to  $C$  might arise that cannot be resolved by promotion after  $S_i$  or by demotion before the action used to achieve  $O:C$ .

These two circumstances are illustrated in Figures 2 and 3. The second of these is a bit tricky. If the threat to  $C$  were resolved by forcing it to come after step  $S_i$  then the



**Figure 2:** Link into a step between repeated open conditions.



**Figure 3:** Threat to  $C$  that can only be resolved by placing  $S_t$  between  $S_k$  and  $S_j$ . The dashed edges indicate ordering constraints, and the thick grey edge indicates a threat.

recursion steps  $S_k$  through  $S_j$  would still be unnecessary. Likewise, if the threat could be resolved by forcing it to come before whatever step is used to establish  $O:C$  the recursion steps  $S_k$  through  $S_j$  would still be unnecessary. Thus, it is only when the threat is forced to come between  $S_k$  and  $S_j$  that this partial plan must be saved.

Our approach for controlling such recursion is two-pronged. First, we use an ordering strategy that prevents expansion of open conditions like  $O:C$  unless/until 1) some other condition links into the loop (as in Figure 2) or 2) a threat appears to the loop condition that must be ordered within the loop (as in Figure 3). Second, we specify conditions where such recursive plans can actually be pruned. To make all of this precise, we first need some definitions:

**Definition 1:** An open condition  $O:C$  is said to be *exactly recursive* if:

1. For every causal link path  $P$  from  $O:C$  to the goal, there is a causal link  $L: S_i \xrightarrow{C} S_j$  in  $P$ .
2. None of the  $S_j$  are ordered with respect to each other.

We refer to the causal links  $L$  as the *root-links* of the recursion. ■

In the example in Figure 1,  $S_j \xrightarrow{C} S_i$  is the only root link for the recursive open condition  $O:C$ .

**Definition 2:** Let  $O:C$  be a recursive open condition, and let  $L$  be its root links. We say that  $D$  is a *loop descendant* for  $O$  if every causal link path from  $D$  to a goal contains a link  $l \in L$ . ■

Note that the condition  $O$  will always be a loop descendant, but there may be many more other (non-recursive) open conditions generated by steps in the loop, and these conditions will generally be loop descendants.

In the example in Figure 1,  $O:C$  is a loop descendant, but any other open conditions resulting from steps from  $S_k$  through  $S_j$  (inclusive) would also be loop descendants.

**Definition 3:** Let  $O:C$  be a recursive open condition, and let  $L$  be the root links of the recursion. A step  $T$  is said to be *loop threat* for  $O$  if:

1.  $T$  plausibly threatens  $C$  ( $T$  has an effect  $\neg C$ ).
2.  $T$  necessarily precedes some  $l \in L$ . ■

In Figure 3, the step  $S_i$  is a loop threat for  $O:C$ .

We can now state our open condition ordering strategy precisely.

**Recursion Ordering Strategy:** Let  $O$  be a recursive open condition. If there is no loop threat for  $O$ , the planner is prevented from working on  $O$  and any other loop descendant  $D$  for  $O$ . ■

In practice, we manage this by marking such open conditions as *suspended*. All suspended conditions for a loop must be re-enabled if either:

1. A link is made to a step between  $O$  and one of its root links,
2. A loop threat for  $O$  appears.<sup>1</sup>

In addition, a specific suspended loop descendent  $D$  will need to be re-enabled whenever a link is made to any step between  $D$  and a root link.

Roughly speaking this strategy suspends all open conditions in a loop, and only re-enables them if either another condition links into the loop or a loop threat occurs.

**Theorem 1:** A partial plan can be pruned if all open conditions in the plan are suspended, and all threats have been resolved.<sup>2</sup> ■

For the example in Figure 1, the open condition  $O:C$  would be suspended along with any other conditions descendant from  $S_j \xrightarrow{C} S_i$ . The plan can be pruned by this theorem if no links (as in Figure 2) are made into the loop, and no loop threats (as in Figure 3) occur.

### 3 Detection and Early Pruning

Although the method described in the previous section can be used as is, there is considerable overhead involved in recognizing recursion. For each new step added to the plan, the planner would need to search each causal link path to the goal, and look to see if any of the newly added open conditions are present on the path. While this is inexpensive for shallow plans, the cost grows significantly with plan depth and complexity. In [13], we introduced a structure called an *operator graph*

<sup>1</sup> This simple strategy still admits some non-minimal plans. In practice, we use a more complex strategy. We do not re-enable the suspended open conditions because it still might be possible to resolve the threat by ordering it before the recursive condition  $O$ . We wait until no other open conditions remain, force the threatening step to come within the loop, and then re-enable the suspended conditions.

<sup>2</sup> In this case, the plan is not minimal (see [7]).

that captures the interaction between operators relevant to a goal and set of initial conditions. In that paper, operator graphs were used for analyzing possible conflicts between operators relevant to a planning problem. However, the operator graph can be used for many other purposes, including recognizing:

1. Open conditions that have the potential to be recursive
2. Open conditions in a partial plan that will never result in links to steps in a suspended loop
3. Open conditions that will never lead to threats to a suspended condition

To explain these further, we first need to review the basics of operator graphs.

An operator graph is a directed bipartite graph containing one operator node for each operator relevant to a goal and one precondition node for each precondition of each such operator. The graph is constructed recursively by working backwards from the goal, according to the following rules:

1. There is an operator node for the Finish operator.
2. If an operator node is in the graph, there is a precondition node in the graph for each precondition of the operator, and a directed edge from the precondition node to the operator node.
3. If a precondition node is in the graph, there is an operator node in the graph for every operator with an effect that unifies with the precondition, and there is a directed edge from the operator node to the precondition node.

Figure 4 shows an operator graph for a simple travel problem where the goal is to be someplace with gas, and there is one operator:


```

Drive(x,y)
  Preconditions:  At(x), Road(x,y)
  Effects        At(y), ¬ At(x)

```

(Lower case letters are used for variables. Constants and relations are capitalized.)

There are two arcs into the goal condition  $At(y)$  indicating that there are two possible ways of achieving it: one by linking to the initial conditions and the other by using the Drive operator. Similarly, there are two arcs into the  $At$  precondition of Drive. It too can potentially be satisfied by binding to the initial conditions, or by using another Drive operation. The other precondition of Drive,  $Road(x,y)$ , can only be satisfied by the initial conditions, but there are several possible ways of doing so, indicated by the bundle of arcs from Start to  $Road(x,y)$ . Likewise, the other goal condition  $Gas(y)$  can only be satisfied by the initial conditions, but again there are several possible ways of doing so.



**Figure 4:** Operator graph for a simple travel problem.

### 3.1 Detecting Recursion

Our first observation is that loops within the operator graph indicate possible recursion during planning. More precisely:

**Theorem 2:** Let  $O$  be an open condition, and let  $P$  be its corresponding precondition node within the operator graph. (That is, the corresponding precondition node for the operator used in the step that gave rise to  $O$ .) If  $O$  is a recursive open condition,  $P$  must be part of a strongly-connected component (SCC)<sup>3</sup> within the operator graph. (The converse is not always true.) ■

The impact of this theorem is that the planner does not have to examine every open condition for possible recursion, only those that correspond to precondition nodes in the operator graph contained in SCCs.

In the graph above, there is one SCC that contains the Drive operator and its precondition  $At(x)$ . As a result, the only open conditions that need to be checked for possible recursion are  $At$  open conditions that result from adding a Drive step to the plan.

SCCs can also be used to limit the amount of search involved in deciding whether a candidate open condition is recursive.

**Theorem 3:** Suppose  $O$  is an open condition that is a possible candidate for recursion, and suppose that the corresponding precondition  $P$  in the operator graph is in a SCC  $K$ . Let  $L_{ij} : S_i \xrightarrow{C} S_j$  be a causal link along a path from  $O$  to the goal. If the operator node in the operator graph corresponding to  $S_j$  is not in the SCC  $K$ , no causal link between  $L_{ij}$  and the goal can be a root of the recursion. ■

The import of Theorem 3 is that the planner can stop looking up the causal link chain as soon as it leaves the SCC. (For the example above, this theorem has no impact, because the first causal link in the chain that is outside the SCC would be the one from Drive to Finish, establishing the goal.)

<sup>3</sup> A strongly connected component is a maximal set of nodes such that there is a directed path in the graph from any node in the set to any other node in the set. Any loop in a directed graph (like the one in Figure 4) will be part of a SCC. SCCs and algorithms for finding them are described in algorithms texts such as [2].

## 3.2 Detecting Potential Links and Threats

It would be nice if we could tell which open conditions in a plan have the potential to link into a suspended loop. Likewise, it would be nice if we could tell which open conditions could potentially lead to loop threats. We can do this by looking at the nodes below the corresponding precondition node in the operator graph.

**Definition 4:** An operator  $R$  is said to be *relevant* to an open condition  $O$  if there is a directed path from  $R$  to the precondition node corresponding to  $O$  in the operator graph. ■


Given a suspension, the set of relevant operators for the remaining open conditions can help us decide whether the plan must be kept or can be discarded.

**Theorem 4:** Let  $P$  be a partial plan with a suspended open condition  $O:C$ . Let  $U$  be the remaining open conditions (including any other suspended open conditions), and let  $R$  be the set of operators relevant to the conditions in  $U$ . The partial plan can be discarded if:

1. No operator in  $R$  appears in the SCC of  $O:C$ .
2. No operator in  $R$  threatens  $C$ .
3.  $P$  contains no unresolved threats to  $C$ . ■

The first of these conditions means that no other operator can link into the suspended loop. The second and third conditions mean that no threats can arise that would force re-enabling of the suspended condition. Every possible completion of this plan will therefore satisfy the conditions of Theorem 1, so the plan can be pruned.

As an example of the application of his theorem, consider the partial plan shown in Figure 5. We've added the goal condition *Rested* and the operator *Sleep* to the problem description. *Sleep* has no preconditions and can be used to achieve *Rested*. Currently, the open condition *At(B)* is suspended because it exactly matches the condition in the causal link  $\text{Drive}(A, B) \xrightarrow{\text{At}(B)} \text{Finish}$ . The relevant operators for *Gas(B)* and *Rested* are *Start* and *Sleep*. Neither of these appear in the SCC for the open condition *At(B)*. As a result, this plan can be pruned.




**Figure 5:** Partial plan with a suspended recursive open condition.



## 4 Instance Recursion

In Sections 2 and 3 we considered only exact recursion. Thus, an open condition  $P(x,A)$  would be recursive if all of its paths to the goal contained a causal link with clause  $P(x,A)$ , but not if the links only contained clauses  $P(y,A)$ ,  $P(x,B)$ , or  $P(x,y)$ . In other words, for exact recursion, the variables and constants in the recursive open condition must be identical to the variables and constants in the root conditions. In the example of Figure 5, the condition  $At(A)$  is not recursive, but the open condition  $At(B)$  is.

While the results of Sections 2 and 3 were sufficient to prune the plan in Figure 5, in general, they are not sufficient to stop all recursion for planning problems in which the operators contain variables. Consider what would happen in our Drive example if the planner did not choose a good order for open condition expansion. The planner could generate an infinite sequence of partial plans like the one shown in Figure 6. While this open condition ordering may look silly, it is exactly the one



**Figure 6:** Recursion resulting from unbound variables.

that would be chosen by such “smart” (and generally effective) strategies as the Least-commitment (LC) strategy used in [10] and Joslin’s LCFR strategy [6]. These strategies choose to expand the open condition with the fewest options. They get into trouble because there are many roads and many gas stations, but only two ways of achieving a condition like  $At(x)$ : adding a Drive step or binding to the single At fact in the initial conditions. If there is no solution to the above problem (because you start on an island and all gas stations are on the mainland), planners with the LC or LCFR strategy would never terminate.

One obvious possibility is to avoid working on conditions such as  $At(x)$  and  $At(y)$  until the variables  $x$  and  $y$  are bound. As it turns out, this isn’t quite right.<sup>4</sup> Instead, we want to avoid working on  $At(x)$  and  $At(y)$  until the variable  $z$  is bound in  $At(z)$ . More precisely, we want to avoid  $At(x)$  and  $At(y)$  as long as they remain an *instance* of the root link,  $At(z)$ .

The reason for this strategy is that until  $z$  is bound, the planner doesn’t know what it is trying to achieve, which, in a recursive situation like this, can lead to lots of

<sup>4</sup> One reason is that all open conditions containing the variables of interest may be recursive, so the variables will never get bound without expanding a recursive open condition. The simplest example of this is a Drive operator that has no Road precondition (the vehicle might be a Humvee or tank). In this case, there would be no other open condition to bind the  $x$  in  $At(x)$ .

irrelevant plans. Another way of thinking about it is this: suppose there is some plan for achieving  $At(x)$ , for some binding of  $x$ . As long as  $At(x)$  remains an instance of  $At(z)$ , that same plan would be directly applicable to  $At(z)$ . Until the planner has reason to believe that  $z=B$  won't work, there's no reason this shorter plan shouldn't be preferred. In the extreme case where  $z$  is never bound (gas is available everywhere), any plan involving recursion would be non-minimal.

As we did for exact recursion, we now provide a formal definition of instance recursion and give a suspension strategy.

**Definition 5:** An open condition  $O:C$  is said to be *instance recursive* if there is a set of variable bindings  $V$  such that:

1. for every causal link path  $P$  from  $O:C'$  to the goal, there is a causal link  $L_{ij} : S_i \xrightarrow{C} S_j$  in  $P$  and  $C' = C|_V$  (where  $C|_V$  refers to  $C$  instantiated with the variable bindings  $V$ ).
2. None of the  $S_j$  are ordered with respect to each other.

As before, we refer to the  $L_{ij}$  as the root links of the recursion. ■

In the above example  $At(x)$  and  $At(y)$  are both instance recursions of the root link  $Drive(y, z) \xrightarrow{At(z)} Finish$ , with  $z$  bound to  $x$  and  $y$  respectively.

**Instance Recursion Strategy:** Let  $O$  be an instance recursive open condition. If there is no loop threat for  $O$  then the planner is prevented from working on  $O$  and any other loop descendant  $D$  for  $O$ . ■

As with exact recursion, we manage this by marking such open conditions as suspended. In this case all suspended conditions for a loop must be re-enabled if either:

1. A link is made into a step between  $O$  and one of its root links.
2. A loop threat for  $O$  appears.
3. Variables bindings are added to the plan so that  $O$  is no longer an instance of one or more of its root links.

In addition a specific suspended loop descendent  $D$  will need to be re-enabled whenever a link is made to any step between  $D$  and a root link.

We can now generalize Theorem 1 to the cases involving instance recursion.

**Theorem 5:** A partial plan can be pruned if all remaining open conditions are suspended (either because of exact or instance recursion) and all threats have been resolved. ■

As before, the argument is that the plan  $P$  is not minimal; if there is a way of achieving a suspended condition  $O:C'$ , that same plan would work for achieving  $O:C$  directly, because  $C'$  is an instance of  $C$ .

As before, we would like to be able to prune such plans earlier. In Section 3 we showed how the operator graph can help recognize when outstanding open conditions can no longer link into or threaten an exact recursion. We can do the same here. However, we also need to make sure that the outstanding open conditions cannot bind any of the variables in the root links. We generalize Theorem 4 as follows.

**Theorem 6:** Let  $P$  be a partial plan with a suspended instance recursive open condition  $O$  and let  $L$  be the set of root links for the recursion. Let  $U$  be the remaining open conditions (including any other suspended open conditions), and let  $R$  be the set of operators relevant to the conditions in  $U$ . The partial plan  $P$  can be pruned if:

1. No operator in  $R$  appears in the SCC of  $O$ .
2. No operator in  $R$  threatens the condition of any root link  $l \in L$ .
3.  $P$  contains no unresolved threats to any root link  $l \in L$ .
4. No open condition in  $U$  contains a variable in any of the root link conditions  $l \in L$ . ■

Although the suspension strategy and pruning theorems for instance recursion look very similar to those for exact recursion, there is a subtle difference. The final condition in Theorem 6 makes it much weaker than Theorem 4, so fewer instance recursive plans can actually be pruned. It is worth noting that instance recursions sometimes turn into exact recursions (e.g. Figure 5) and are then subject to the stronger Theorem 4.

## 5 Experiments

The exact and instance suspension and pruning strategies were implemented in our ONLP planner (Operator Graph-based partial order planner). We tested the recursion algorithms on a number of problems drawn from several domains including: Russell's Tire World domain [10], Bulldozer [10], a railroad meet/pass planning domain, towers of Hanoi, Monkey and Bananas, Weld's Refrigerator domain, a noncombatant evacuation operation (NEO) domain, and blocks world. Many of the descriptions for these domains were adapted from the domain library in UCPOP [9].


ONLP implements a number of strategies for operator graph and plan analysis. For our tests, we used the following settings:

- Open Condition Ordering: Least Commitment [10] – select the open condition that can be satisfied in the fewest ways.
- Threat Delay Strategy: DMIN [14] – a least commitment threat resolution strategy similar to DUNF [10].

- Variable Analysis: substitute the binding for each variable in the operator schemas if there is only one possible binding for that variable.
- Cost Function:  
 $\# \text{ Steps} + \# \text{ Unsuspended Open Conds} + K (\# \text{ Suspended Open Conds})$   
 The constant  $K$  penalizes partial plans that have suspended open conditions. The values used in our tests were  $K = 1$  (no penalty),  $K = 2$  and  $K = 4$ .
- Search Algorithm: Stable best-first--A best first search algorithm that always breaks ties in the order that the plans were added to the search queue (LIFO).

The results of these tests are shown below. In these charts,  $N_{\text{normal}}$  and  $N_{\text{suspend}}$  denote the number of plans generated by the planner with and without recursion suspension and pruning. (These counts do *not* include the number of plans explored during threat resolution, only the plans generated using add-link or add-step. See [10] and [14] for explanation.)

Figure 7 plots the reduction in search space size realized by using recursion suspension and pruning ( $N_{\text{normal}}/N_{\text{suspend}}$ ) against the size of the search space when no suspension and pruning is used ( $N_{\text{normal}}$ ). Two data series are plotted in this chart: one for  $K=1$  and one for  $K=4$ .




**Figure 7:** Search space improvement when using recursion suspension and pruning.

Almost all of the points in Figure 7 are above the X-Axis, which indicates that there was some degree of search space reduction achieved for most problems. The improvement ranges from none to as much as 139 times on the tire changing problem. Generally, the improvement is greater for the larger domains and harder problems. One reason is that these plans tend to be longer, and the odds of a plan being non-minimal increases with the length of the plan (more opportunities for recursion). Note that better performance is obtained using the penalty  $K=4$  for

suspended open conditions. This reflects the high probability that the best plan does not contain recursion.

Figure 8 plots the improvement in search time using recursion suspension and pruning ( $T_{\text{normal}}/T_{\text{suspend}}$ ) against the time (in seconds) required for solving each problem with no suspension and pruning ( $T_{\text{normal}}$ ).



**Figure 8:** Time improvement when using recursion suspension and pruning

As with search space size, all points are either close to, or above the X-axis. This indicates that the overhead of recursion checking and suspension is sufficiently small that it has little negative impact on planner performance, even when no search space reduction is realized by the strategy. The overhead of building the operator graph, and finding strongly connected components was negligible. For small problems it was typically less than 0.1 seconds, and for all problems it was less than 1 second.

## 6 Conclusions

The recursion suspension and pruning strategy outlined in Sections 2, 3, and 4 significantly reduce search time and space for many problems and, in particular, they reduce search time significantly for large recursive problems. Using the operator graph to help detect recursion keeps the overhead associated with this technique to a minimum. Finally, our test results indicate that penalizing plans with large numbers of suspended open conditions can further improve performance.

## Acknowledgments

We first implemented the suspension method in early 1993. Dan Weld encouraged us to finally write it up and provided comments on an early draft. Thanks to Tony Barrett for collecting and making many of the test domains available. This work was supported by DARPA contract F30602-91-C-0031 and by Rockwell.

## References

1. Barrett, A. and Weld, D., Partial order planning: evaluating possible efficiency gains. *Artificial Intelligence*, vol 67(1), pages 71–112, 1994.
2. Cormen, T., Leiserson, C., and Rivest, R., *Introduction to Algorithms*. McGraw Hill, 1991.
3. Feldman, R. and Morris, P., Admissible criteria for loop control in planning. In *Proc. AAAI-90*, pages 151–157, 1990.
4. Etzioni, O., Acquiring search-control knowledge via static analysis. *Artificial Intelligence*, vol 62(2), pages 255–301, 1993.
5. Fikes, R., Hart, P., and Nilsson, N., Learning and executing generalized robot plans. *Artificial Intelligence*, vol 3(4), pages 251–288, 1972.
6. Joslin, D. and Pollack, M., Least cost flaw repair: a plan refinement strategy for partial-order planning. In *Proc. AAAI-94*, pages 1004–1009, 1994.
7. Kambhampati, S., Admissible pruning strategies based on plan minimality for plan-space planning. In *Proc. IJCAI-95*, pages 1627 – 1633, 1995.
8. McAllester, D. and Rosenblitt, D., Systematic nonlinear planning. In *Proc. AAAI-91*, pages 634–639, Anaheim, CA, 1991.
9. Penberthy, J. and Weld, D., UCPOP: A sound, complete, partial order planner for ADL. In *Proc. KR-92*, pages 189–197. 1992.
10. Peot, M. and Smith, D., Threat removal strategies for partial-order planning. In *Proc. AAAI-93*, pages 492–499, 1993.
11. Rich, E. and Knight, K., *Artificial Intelligence*. Second edition, McGraw Hill, 1991.
12. Smith, D., Genesereth, M., and Ginsberg, M., Controlling recursive inference. *Artificial Intelligence*, vol 30(3), pages 343–389, 1986.
13. Smith, D. and Peot, M., Postponing threats in partial order planning. In *Proc. AAAI-93*, pages 500–506, 1993.
14. Smith, D. and Peot, M., *A Note on the DMIN strategy*. Technical Memo, Rockwell Palo Alto Laboratory, 1993. Available on-line from <http://www.rpal.rockwell.com:80/~de2smith/publications.html>
15. Tate, A., Generating project networks. In *Proc. IJCAI-77*, pages 888-893, 1977.